

I Claim:

1. An apparatus for forwarding packets between ports, said ports associated with hosts having one or more of a layer 2 address, a layer 3 address and a socket number, said apparatus

5 comprising:

a flow table having a plurality of records;

a switch module coupled to said ports and said flow table, said switch module

comprising:

means for forwarding a first packet between a first host and a second host in accordance with a first record in said flow table and said layer 2 addresses of said first and second hosts, and

means for forwarding a second packet between a third host and a fourth host in accordance with a second record in said flow table and said layer 3 addresses of said third and fourth hosts.

2. An apparatus according to claim 1, wherein said flow table comprises:

an address resolution hash table comprising at least first and second hash entries, said first hash entry corresponding to said layer 2 address of one of said first and second hosts, said second hash entry corresponding to said layer 3 address of one of said third and fourth hosts; and

an address resolution records table, said first and second records being stored in said

address resolution records table and linked to said first and second hash entries in said address resolution hash table, said switch module having means for retrieving said first and second records using portions of said layer 2 address of said one of said first and second hosts and said layer 3 address of said one of said third and fourth hosts, respectively.

5 3. An apparatus according to claim 2, wherein said means for retrieving further uses said socket number of said one of said third and fourth hosts for retrieving said second record.

10 4. An apparatus according to claim 1, wherein said first packet is associated with a non-IP/IPX protocol, and said second packet is associated with an IP/IPX protocol.

5. An apparatus according to claim 1, further comprising:

15 a CPU coupled to said flow table and communicating with said switch module, said switch module further comprising means for alerting said CPU when a third packet between a fifth host and a sixth host arrives at one of said ports and a third record containing forwarding information between said fifth host and said sixth host does not exist in said flow table, said CPU creating said third record and storing said third record in said flow table, said switch module thereafter forwarding said third packet between said fifth host and said sixth host in accordance with said third record and said addresses of said fifth and sixth hosts.

20

6. An apparatus according to claim 5, wherein said switch module and said flow table are together comprised of an ASIC.

7. An apparatus according to claim 2, further comprising:

5 a CPU coupled to said flow table and communicating with said switch module, said switch module further comprising:

10 means for alerting said CPU when a third packet between a fifth host and a sixth host arrives at one of said ports and a third record containing forwarding information between said fifth host and said sixth host does not exist in said flow table, said CPU creating said third record and storing said third record in said address resolution record table, said CPU further creating a hash entry in said address resolution hash table that is linked to said third record and is associated with a portion of said address of one of said fifth and sixth hosts,

15 means for hashing onto said address resolution record table in accordance with said portion of said address of said one of said fifth and sixth hosts,

20 means for retrieving said third record based on said hash, said switch module thereafter forwarding said third packet between said fifth host and said sixth host in accordance with said third record and said addresses of said fifth and sixth hosts.

8. An apparatus according to claim 1, wherein said first and second hosts belong to different networks.

5 9. An apparatus according to claim 1, wherein said third and fourth hosts belong to different networks.

10. An apparatus according to claim 1, wherein said first, second, third and fourth hosts belong to the same network.

10 11. An apparatus for forwarding packets between ports, said ports including a first port associated with a first host having a first layer 2 address and a first layer 3 address, and a second port associated with a second host having a second layer 2 address and a second layer 3 address, a layer 2 flow of packets between said first and second hosts being based on said first and second layer 2 addresses, a layer 3 flow of packets between said first and second hosts being based on said first and second layer 3 addresses, said apparatus comprising:

15 a flow table comprising a plurality of address resolution records including:

a first address resolution record that corresponds said first port with said first layer 2 address,

a second address resolution record that corresponds said first port with said first layer 3 address,

a third address resolution record that corresponds said second port with said second layer 2 address, and

a fourth address resolution record that corresponds said second port with said second layer 3 address; and

5 a switch module coupled to said ports and said flow table that detects said layer 2 flow arriving at one of said first and second ports, and forwards packets belonging thereto to the other of said first and second ports based on said first and third address resolution records, said switch module also detects said layer 3 flow arriving at one of said first and second ports, and forwards packets belonging thereto to the other of said first and second ports based on said second and fourth address resolution records.

10 12. An apparatus as defined in claim 11, further comprising:

a CPU coupled to said flow table and communicating with said switch module, said CPU creating said first, second, third and fourth address resolution records in said address resolution record table.

15 13. An apparatus as defined in claim 12, wherein said switch module includes:

Sub
a 17 a switch engine that determines whether said first and third address resolution records exist in said address resolution record table when said first flow is detected, and whether
20 said second and fourth address resolution records exist in said address resolution record table when

said second flow is detected; and

a CPU interface that sends a first message to said CPU when said first flow is detected and said switch engine determines that said first and third address resolution records do not exist in said address resolution record table, said CPU interface sends a second message to said CPU when said second flow is detected and said switch engine determines that said second and fourth address resolution records do not exist in said address resolution record table, said CPU creating said first and third address resolution records in response to said first message, and said second and fourth address resolution records in response to said second message.

10 14. An apparatus as defined in claim 11, wherein said layer 3 flow is in accordance with one of IP and IPX protocols.

Sub
a2 } 15. An apparatus as defined in claim 14, wherein said layer 2 flow is in accordance with all protocols except IP and IPX..

15 16. An apparatus as defined in claim 11, wherein said flow table further includes an address resolution hash table comprising a plurality of address resolution hash records including:

a first address resolution hash record that corresponds a portion of said first layer 2 address with said first address resolution record;

20 a second address resolution hash record that corresponds a portion of said first layer

3 address with said first address resolution record;

a third address resolution hash record that corresponds a portion of said second layer

2 address with said third address resolution record; and

a fourth address resolution hash record that corresponds a portion of said second

5 layer 3 address with said fourth address resolution record.

17. An apparatus as defined in claim 13, wherein said flow table further includes an address resolution hash table comprising a plurality of address resolution hash records including:

a first address resolution hash record that corresponds a portion of said first layer 2

10 address with said first address resolution record;

a second address resolution hash record that corresponds a portion of said first layer

3 address with said first address resolution record;

a third address resolution hash record that corresponds a portion of said second layer

2 address with said third address resolution record; and

15 a fourth address resolution hash record that corresponds a portion of said second

layer 3 address with said fourth address resolution record,

said switch module accessing said first, second, third and fourth address resolution records in accordance with said corresponding portions of said addresses and said first, second, third and fourth address resolution record hash records, respectively,

20 said CPU linking said first and third address resolution hash records with said first

and third address resolution records, respectively, in response to said first message, and linking said second and fourth address resolution hash records with said second and fourth address resolution records, respectively, in response to said second message.

5 18. An apparatus as defined in claim 11, wherein said first and second hosts have first and second sockets, respectively, said layer 3 flow being further based on said first and second sockets, said second address resolution record further corresponding said first port with said first socket, said fourth address resolution record further corresponding said second port with said second socket.

10 19. An apparatus as defined in claim 18, further comprising:
a CPU coupled to said flow table and communicating with said switch module, said CPU creating said first, second, third and fourth address resolution records in said address resolution record table.

15 20. An apparatus as defined in claim 19, wherein said switch module includes:

a switch engine that determines whether said first and third address resolution records exist in said address resolution record table when said first flow is detected, and whether said second and fourth address resolution records exist in said address resolution record table when
20 said second flow is detected; and

a CPU interface that sends a first message to said CPU when said first flow is detected and said switch engine determines that said first and third address resolution records do not exist in said address resolution record table, said CPU interface sends a second message to said CPU when said second flow is detected and said switch engine determines that said second and fourth address resolution records do not exist in said address resolution record table, said CPU
5 creating said first and third address resolution records in response to said first message, and said second and fourth address resolution records in response to said second message.

21. An apparatus as defined in claim 20, wherein said flow table further includes an address resolution hash table comprising a plurality of address resolution hash records including:
10

a first address resolution hash record that corresponds a portion of said first layer 2 address with said first address resolution record;

a second address resolution hash record that corresponds a portion of said first layer
3 address with said first address resolution record;

a third address resolution hash record that corresponds a portion of said second layer
15 2 address with said third address resolution record; and

a fourth address resolution hash record that corresponds a portion of said second layer 3 address with said fourth address resolution record,

said switch module accessing said first, second, third and fourth address resolution
20 records in accordance with said corresponding portions of said addresses and said first, second,

third and fourth address resolution record hash records, respectively,

said CPU linking said first and third address resolution hash records with said first and third address resolution records, respectively, in response to said first message, and linking said second and fourth address resolution hash records with said second and fourth address resolution records, respectively, in response to said second message.

22. An apparatus as defined in claim 11, wherein said switch module is comprised of an ASIC.

23. An apparatus as defined in claim 11, wherein said switch module and said flow table are together comprised of an ASIC.

24. An apparatus as defined in claim 11, wherein said first and second hosts belong to different networks.

25. A method of forwarding packets between ports of a switch, said ports being associated with hosts having addresses, said method comprising.

identifying a first address of a first one of said hosts;

identifying a first port associated with said first host;

creating a record that corresponds said first address with said first port;

storing said record in a table;

linking said record to a hash;

associating said hash with a portion of said first address;

receiving a data packet at a second port;

extracting said first address from said data packet;

retrieving said record by hashing onto said table using said portion of said first

address; and

forwarding said data packet to said first port in accordance with said record.

26. A method as defined in claim 25, further comprising:

corresponding a swap address with said record;

receiving a second data packet at a third port;

routing said second data packet by forwarding said second data packet to said third
port in accordance with said record and swapping an address within said second data packet with
said swap address in accordance with said record.

27. A method as defined in claim 25, further comprising:

identifying a second address of a second one of said hosts;

identifying a second port associated with said second host;

creating a second record that corresponds said second address with said second port;

storing said second record in said table;

linking said record to said hash;

associating said hash with a portion of said second address;

sorting said first and second records in said table based on said first and second

5 addresses;

receiving a second data packet at a second port;

extracting said second address from said second data packet;

retrieving said second record by hashing onto said table using said portion of said

second address and searching among said first and second records based on said second address;

10 and

forwarding said data packet to said second port in accordance with said second

record.

28. A method as defined in claim 25, wherein said second port is associated with a second

15 one of said hosts, said first and second hosts belonging to different networks.

SUB
Q5 29. A method of forwarding packets between ports of a switch, said ports being associated
with hosts having one or more of an Ethernet address, an IP address, an IPX address, and a socket

number, said method comprising:

20

preparing a flow table comprising a plurality of records;

receiving a first packet at a first port;

extracting a protocol identifier from said first packet;

if said protocol identifier is IP or IPX:

determining whether said first packet needs to be switched within the
same network or routed between different networks;

if said first packet needs to be switched, switching said first packet
between said first port and a second port based on a first record from said
flow table and one of said IP or IPX address of said first packet; and

if said first packet needs to be routed, routing said first packet
between said first port and a third port based on a second record from said
flow table and one of said IP or IPX address of said first packet; and

if said protocol identifier is not IP or IPX:

switching said first packet between said first port and a third port
based on a third record from said flow table and said Ethernet address of
said first packet.

30. A method as defined in claim 29, wherein said switching and routing steps each
comprise:

determining whether a record associated with an address of said first packet exists in

20 said flow table;

if said record does not exist, alerting a CPU to create said record in said flow table;
forwarding said first packet to a port designated in said record.

5 31. A method as defined in claim 29, wherein if said protocol identifier is IP or IPX and said packet needs to be switched, said step of switching said first packet between said first port and said second port is further based on said socket number of said first packet.

10 32. A method of forwarding packets between ports of a switch, said ports including a first port associated with a first host having a first layer 2 address and a first layer 3 address, and a second port associated with a second host having a second layer 2 address and a second layer 3 address, a layer 2 flow of packets between said first and second hosts being based on said first and second layer 2 addresses, a layer 3 flow of packets between said first and second hosts being based on said first and second layer 3 addresses, said method comprising:

15 preparing a first address resolution record in a flow table that corresponds said first port with said first layer 2 address;

preparing a second address resolution record in said flow table that corresponds said first port with said first layer 3 address;

preparing a third address resolution record in said flow table that corresponds said second port with said second layer 2 address ;

20 preparing a fourth address resolution record in said flow table that corresponds said

second port with said second layer 3 address;

detecting said layer 2 flow arriving at one of said first and second ports and forwarding packets belonging thereto to the other of said first and second ports based on said first and third address resolution records; and

5 detecting said layer 3 flow arriving at one of said first and second ports, and forwarding packets belonging thereto to the other of said first and second ports based on said second and fourth address resolution records.

33. A method as defined in claim 32, further comprising:

10 determining whether said first and third address resolution records exist in said flow table when said first flow is detected;

determining whether said second and fourth address resolution records exist in said flow table when said second flow is detected;

15 creating said first and third address resolution records when said first flow is detected and when it is determined that said first and third address resolution records do not exist in said flow table; and

creating said second and fourth address resolution records when said second flow is detected and when it is determined that said second and fourth address resolution records do not exist in said flow table.

20

34. A method as defined in claim 32, wherein said layer 3 flow is in accordance with one of IP and IPX protocols.

sub
ale
5 35. A method as defined in claim 34, wherein said layer 2 flow is in accordance with all protocols except IP and IPX..

36. A method as defined in claim 32, further comprising:

preparing a first address resolution hash record in said flow table that corresponds a portion of said first layer 2 address with said first address resolution record;

10 preparing a second address resolution hash record in said flow table that corresponds a portion of said first layer 3 address with said first address resolution record;

preparing a third address resolution hash record in said flow table that corresponds a portion of said second layer 2 address with said third address resolution record; and

15 preparing a fourth address resolution hash record in said flow table that corresponds a portion of said second layer 3 address with said fourth address resolution record.

37. A method as defined in claim 33, further comprising:

preparing a first address resolution hash record in said flow table that corresponds a portion of said first layer 2 address with said first address resolution record;

20 preparing a second address resolution hash record in said flow table that

corresponds a portion of said first layer 3 address with said first address resolution record;

preparing a third address resolution hash record in said flow table that corresponds a portion of said second layer 2 address with said third address resolution record;

5 preparing a fourth address resolution hash record in said flow table that corresponds a portion of said second layer 3 address with said fourth address resolution record;

linking said first and third address resolution hash records with said first and third address resolution records, respectively, when said first flow is detected and when it is determined that said first and third address resolution records do not exist in said flow table in response to said first message; and

10 linking said second and fourth address resolution hash records with said second and fourth address resolution records, respectively, when said second flow is detected and when it is determined that said second and fourth address resolution records do not exist in said flow table.

38. A method as defined in claim 37, wherein said step of forwarding said packets
15 belonging to said first flow includes accessing said first and third address resolution records in accordance with said corresponding portions of said addresses and said first and third address resolution record hash records, respectively, and

wherein said step of forwarding said packets belonging to said second flow includes
accessing said second and fourth address resolution records in accordance with said corresponding
20 portions of said addresses and said second and fourth address resolution record hash records,

respectively.

39. A method as defined in claim 32, wherein said first and second hosts have first and second sockets, respectively, said layer 3 flow being further based on said first and second sockets, said second address resolution record further corresponding said first port with said first socket, said fourth address resolution record further corresponding said second port with said second socket.

40. A method as defined in claim 39, further comprising:

determining whether said first and third address resolution records exist in said flow table when said first flow is detected;

determining whether said second and fourth address resolution records exist in said flow table when said second flow is detected;

creating said first and third address resolution records when said first flow is detected and when it is determined that said first and third address resolution records do not exist in said flow table; and

creating said second and fourth address resolution records when said second flow is detected and when it is determined that said second and fourth address resolution records do not exist in said flow table.

41. A method as defined in claim 40, further comprising:

preparing a first address resolution hash record in said flow table that corresponds a portion of said first layer 2 address with said first address resolution record;

5 preparing a second address resolution hash record in said flow table that corresponds a portion of said first layer 3 address with said first address resolution record;

preparing a third address resolution hash record in said flow table that corresponds a portion of said second layer 2 address with said third address resolution record;

preparing a fourth address resolution hash record in said flow table that corresponds a portion of said second layer 3 address with said fourth address resolution record;

10 linking said first and third address resolution hash records with said first and third address resolution records, respectively, when said first flow is detected and when it is determined that said first and third address resolution records do not exist in said flow table in response to said first message; and

15 linking said second and fourth address resolution hash records with said second and fourth address resolution records, respectively, when said second flow is detected and when it is determined that said second and fourth address resolution records do not exist in said flow table.

42. A method as defined in claim 41, wherein said step of forwarding said packets belonging to said first flow includes accessing said first and third address resolution records in accordance with said corresponding portions of said addresses and said first and third address

20

resolution record hash records, respectively, and

wherein said step of forwarding said packets belonging to said second flow includes
accessing said second and fourth address resolution records in accordance with said corresponding
portions of said addresses and said second and fourth address resolution record hash records,
5 respectively.

43. A method as defined in claim 32, wherein said first and second hosts belong to different
networks.

add
a7
add B3